Appendix A. Radiometric Processing Process Flow

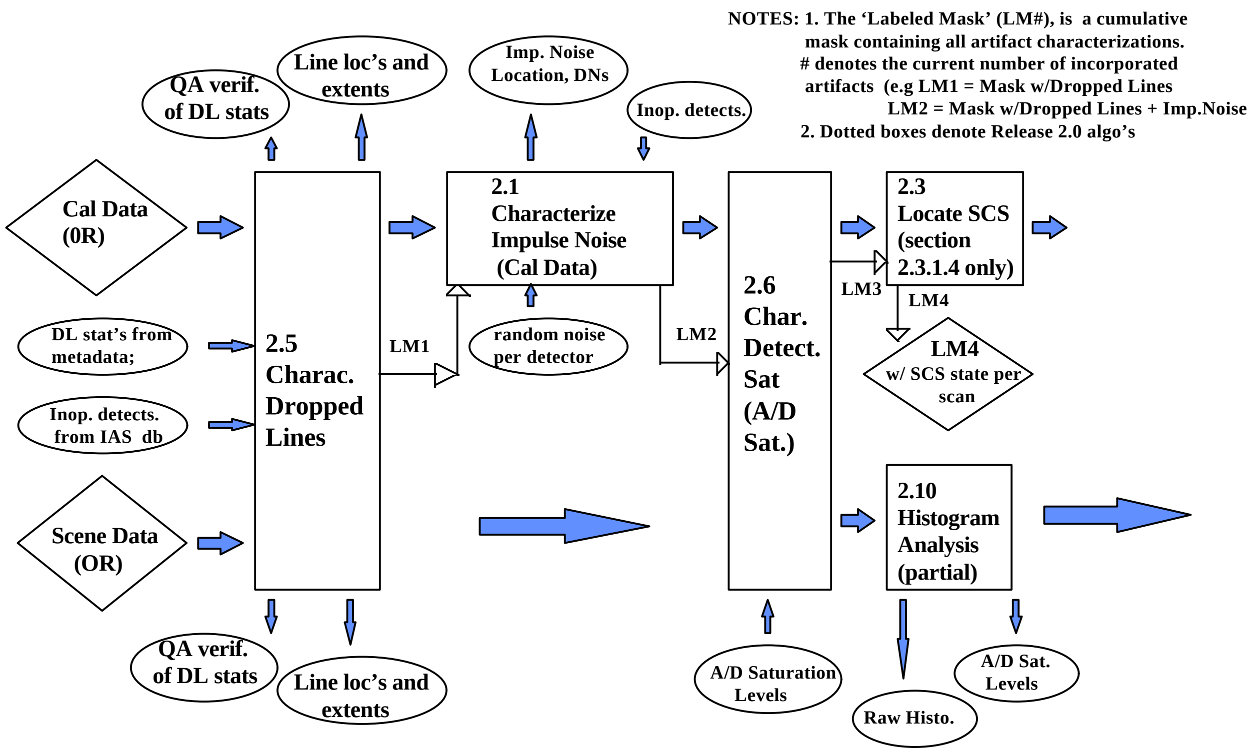

Level 1R Processing (Day Scene)
Release 1.0
Step 1. 0R Radiometric Characterization

(update 3/4/97)

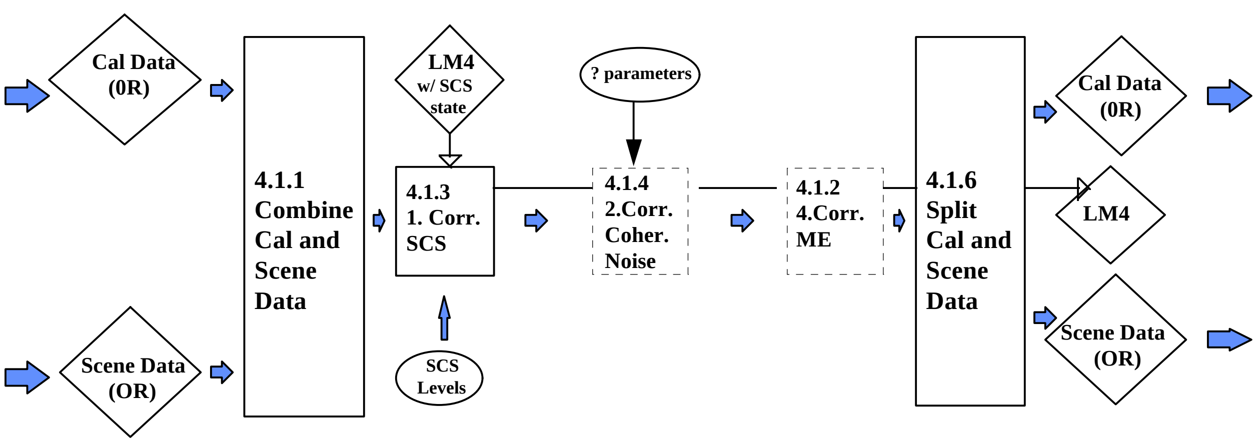

Note: Dotted boxes denote Release 2.0 algorithms

Level 1R Processing (Day Scene)
Release 1.0
Step 2. Pre-1R Correction
(revised 3/4/97)

Notes: 1. Dotted boxes denote Release 2.0 algo's
2. For algorithm "3.4 Process IC", see detailed flow in Step 3a.

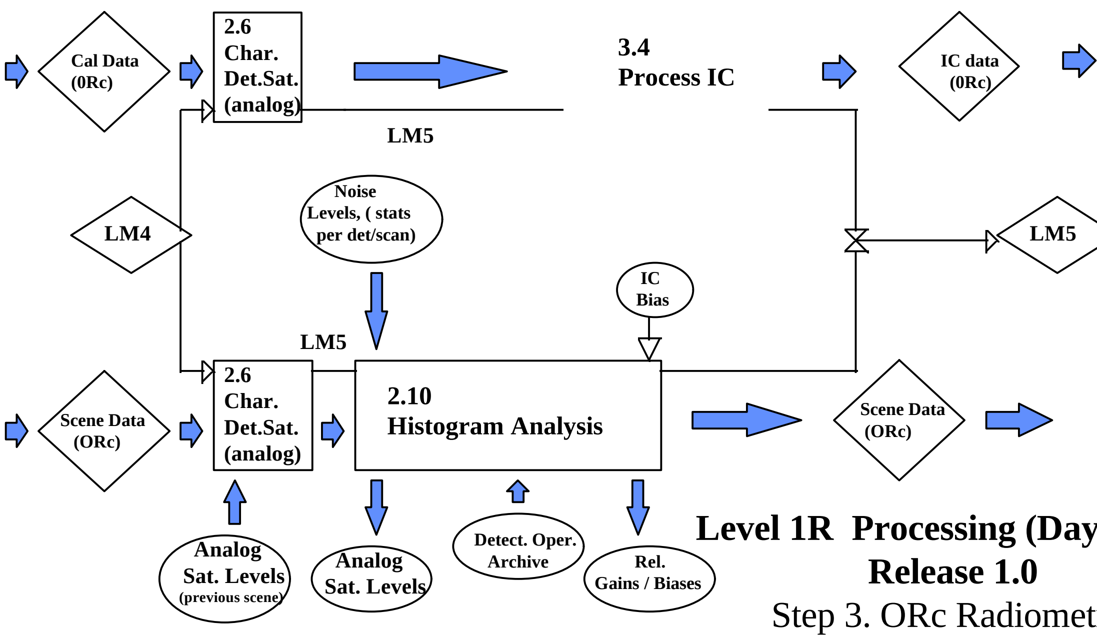

Level 1R Processing (Day Scene)
Release 1.0
Step 3. ORc Radiometric Characterization/Calibration
(revised 3/4/97)

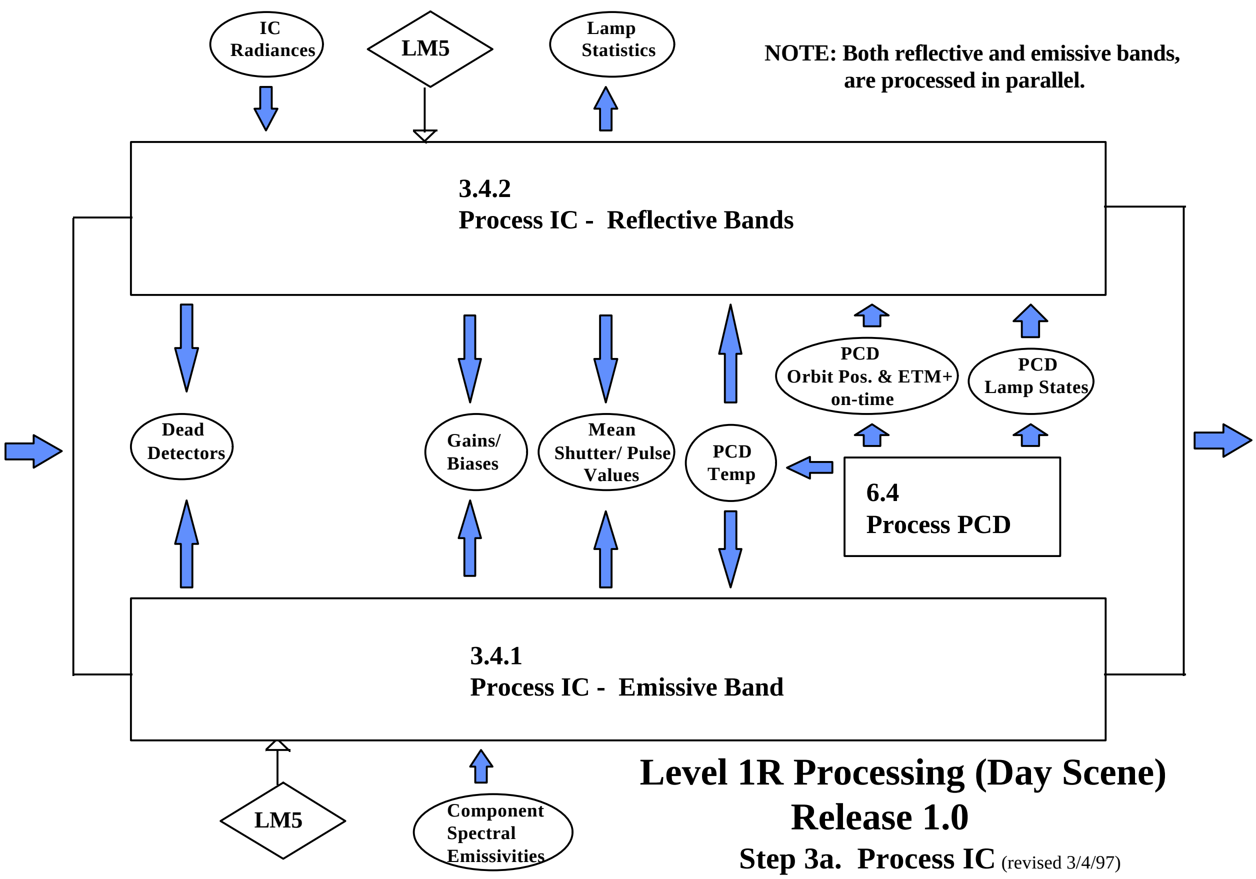

Level 1R Processing (Day Scene)
Release 1.0
Step 3a. Process IC (revised 3/4/97)

Note: Dotted boxes and lines denote Release 2.0 algo's and I/O gains

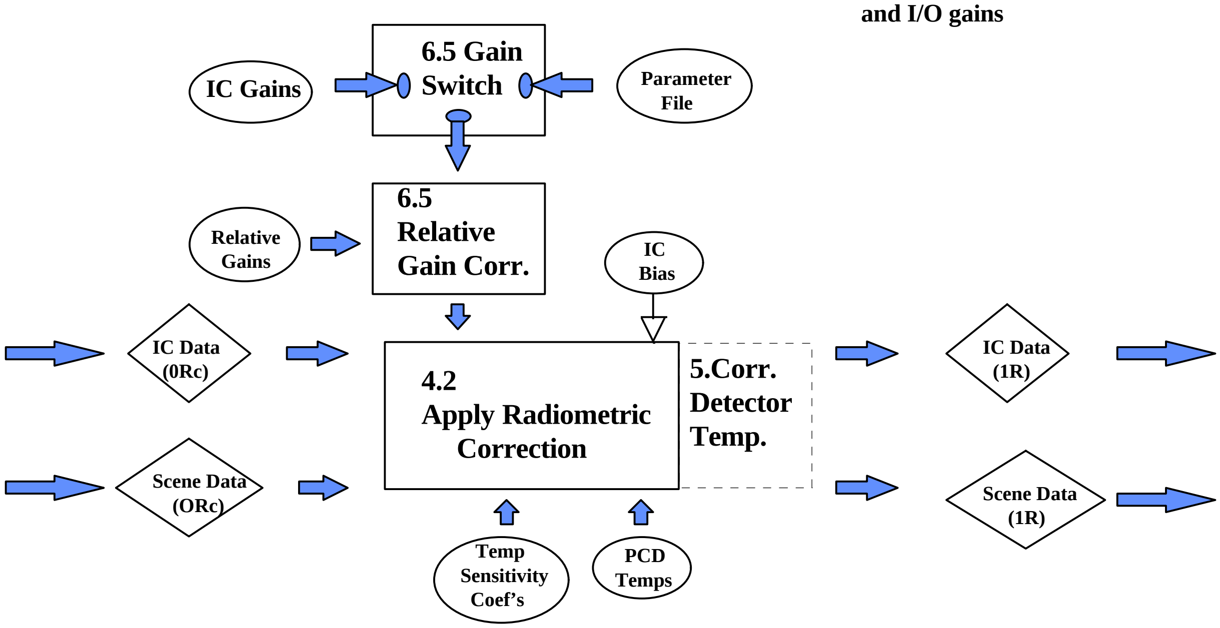

Level 1R Processing (Day Scene)
Release 1.0

Step 4. 1R Correction

(revised 3/4/97)

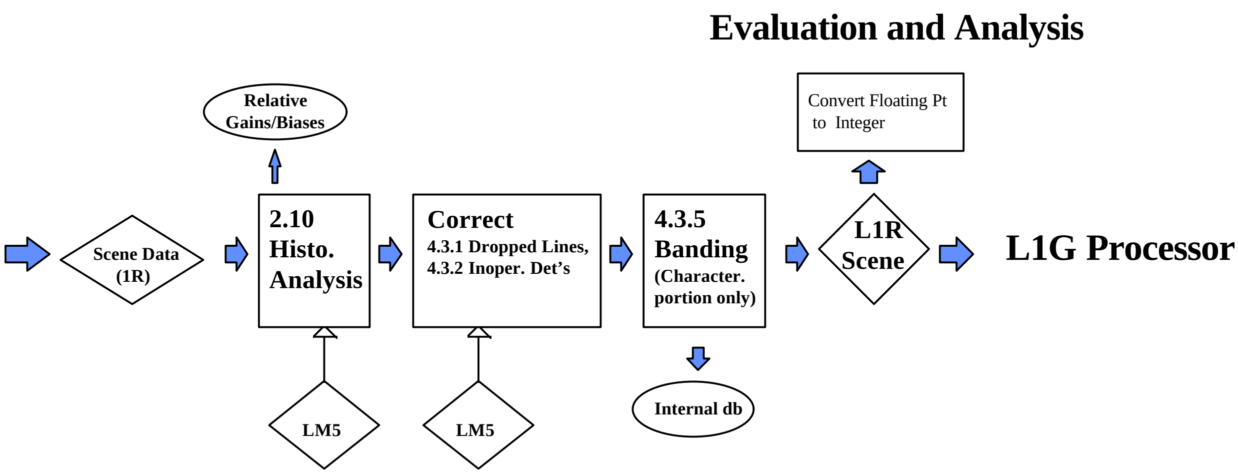

Level 1R Processing (Day Scene)
Release 1.0

Step 5.0 1R Radiometric Characterization/Correction

(Scenario 1: No Correction for Striping and Banding Effects)

(revised 3/4/97)

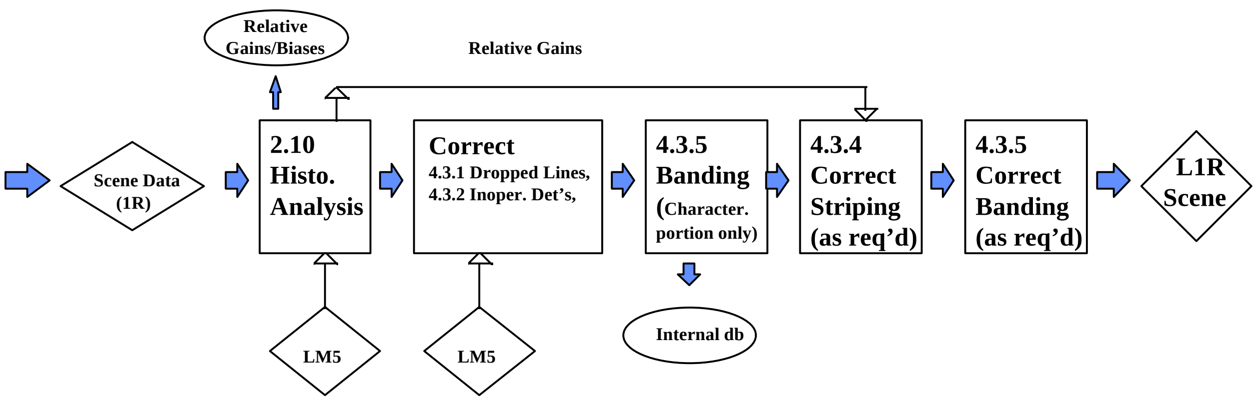

Level 1R Processing (Day Scene)
Release 1.0

Step 5.0 1R Radiometric Characterization/Correction

(Scenario 2: Correct for Striping and Banding Effects IFF Necessary)

(revised 3/4/97)

Appendix B. QASE Simulation Model

Although the spreadsheet model provides accurate mean utilizations and mean service times, it cannot provide information about the dynamic character of the LPGS. In particular, the CPU, FDDI, RAID, and other system resources may become overutilized for protracted times during the normal daily cycle of operations, even though the average values for these components appear to be reasonable. To examine this possibility, a discrete-event simulation model was built.

B.1 Problem Statement

For purposes of this analysis, LPGS performs three types of operations on scene data: nominal processing, reprocessing, and anomaly analysis processing. Nominal processing executes the following activities serially:

1. Data ingestion
2. Radiometric processing
3. Geometric processing
4. Product formatting
5. Product transfer

All bands are processed at each stage before moving to the next stage. The baseline case assumes that 25 scenes receive nominal processing over a 24-hour period, and that this processing can be started uniformly over that planning horizon.

It is assumed that ten percent of the cases will require reprocessing. Reprocessing requires resources at the same levels as the radiometric processing, geometric processing, and product formatting portions of the nominal processing.

In addition, it is assumed that quality assurance checks detect a potential problem with the products for a subset of scenes needs to be reprocessed. When that happens, anomaly analysis executes a set of procedures to identify and correct the problem. Once the problem has been diagnosed, the scene is rerun with the new parameters before it is distributed to LPGS customers. These steps can include the following:

1. Running a set of LPGS benchmarks
2. Executing diagnostics on the offending scene
3. Reprocessing the data with corrected information
4. Transferring the LPGS products.

Benchmarking, diagnostics, and reprocessing each require resources at the same levels as the radiometric processing, geometric processing, and product formatting portions of the nominal processing, since much of the same software is exercised on data sets of equivalent sizes. The number of scenes requires anomaly analysis varies depending on the workload scenarios. It is assumed that three scenes (i.e. 100 percent of scenes to be reprocessed) will require anomaly analysis for the baseline workload scenario.

To complicate matters, management plans to have staff present only during the prime shift. Since anomaly analysis requires manual inspection of the data, this means that all anomalous cases need to be examined during that time.

One strategy is to begin the analysis of a day's worth of cases at the beginning of each prime shift. A second strategy is to have an analyst present at each shift, and to analyze the cases uniformly at the rate of one anomaly analysis thread activation per shift.

To assess the effects of these strategies, management wants to compare the loadings and turnaround times for nominal processing and anomaly analysis over a five-day period. Utilizations are to be computed both in the long-term sense of the spreadsheet and over ten-minute sampling intervals.

Service time measurements include the smallest, largest, and average service times over the five-day sampling period. The number of nominal and anomaly thread activations and completions is also desired.

In the baseline case, the total number of arrivals for each day is set at 25. For the nominal case, this means that there are 25 scenes that begin nominal processing, three scenes that require reprocessing, and no cases that require anomaly analysis. In the case of anomaly analysis, there are three scenes that require anomaly analysis (and reprocessing). Either all three scenes start at once each day (at the same time), or the scenes are started uniformly over a 24-hour period.

In addition to the baseline case, LPGS management is considering increasing capacity of the system to accommodate 100 scenes per day. Ten percent of the cases (ten scenes) require reprocessing. Among the ten scenes to be reprocessed, six scenes require anomaly analysis. The number of CPUs increases from four to 16, with a 14 percent degradation in performance due to increased overhead. The number of RAID groups is increased from one to four with four RAIDs in each group.

B.2 Simulation Model

To assess the dynamic characteristics of the LPGS, a simulation model of the system using the QASE RT system performance modeling package was constructed. QASE RT, sold by Advanced System Technology, Inc. (AST), was originally developed under the auspices of Code 510. The Nominal scenario, shown in Figure B–1, ingests WRS scenes from the ECS, performs radiometric and geometric processing, and transmits the Level 1 products back to the ECS.

Figure B–2 shows anomaly analysis processing. Scenes which exhibit anomalous characteristics must be examined and reprocessed. Because the data must be examined pixel-by-pixel, both nominal and anomalous processing are computing-intensive operations.

LPGS system performance was analyzed by simulating the effect of the workload on the system hardware and software architecture. The LPGS workload is defined by the times between the processing of scenes was initiated. Only a critical subset of the hardware was modeled.

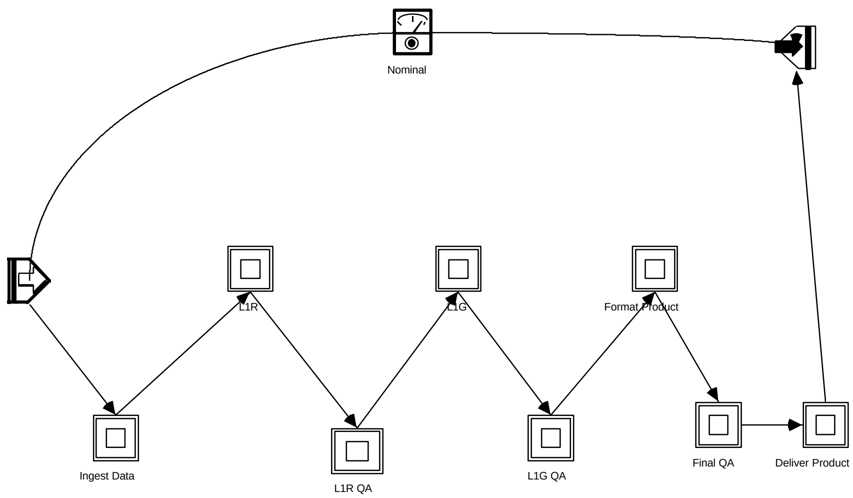

Figure B–1. LPGS Nominal Processing Flow

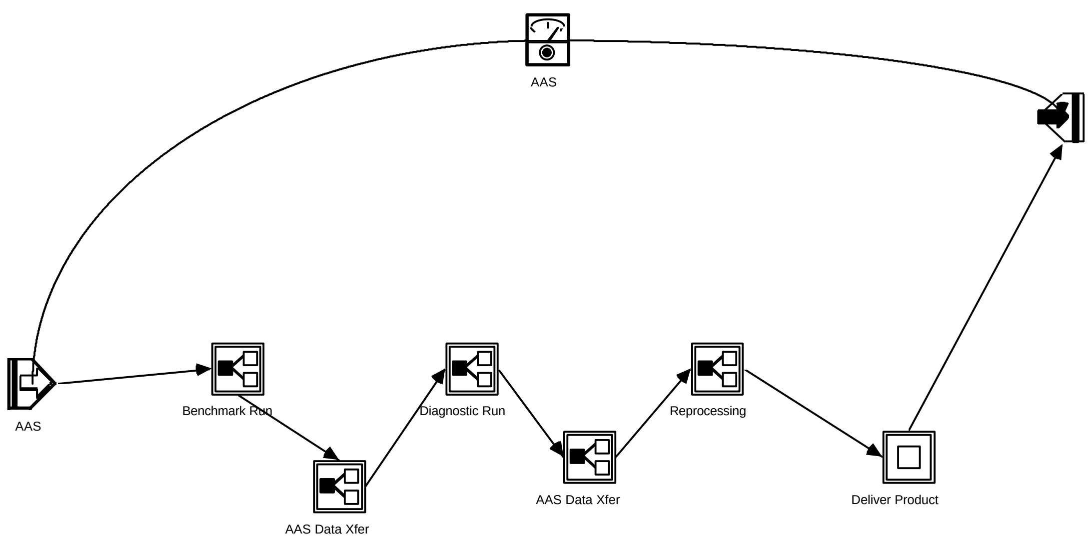

Figure B–2. LPGS Anomaly Processing Flow

B.3 Results

The simulation model provides both summary and dynamic results.

B.3.1 Summary Results

Table B–1 summarizes the response times by workload scenario for a daily workload of processing 28 scenes (including reprocessing three scenes). Table B–2 summarizes the response times by workload scenario for a daily workload of processing 110 scenes (including reprocessing ten scenes). The service time column shows the expected service time for the thread without any queuing effects.

Table B–1. Comparative Response Times (workload = 25 + 3 scenes/day)

Comparative Response Times by Type of Processing				
Response Time	**Service Time**	**Nominal Nom=25/Day Reproc=3/Day AAS = 0/Day**	**Batch Nom=25/Day Reproc=3/Day AAS = 3/Day**	**Uniform Nom=25/Day Reproc=3/Day AAS = 3/Day**
Nominal				
Minimum	N/A	97.68 min	96.60 min	96.61 min
Mean	96.61 min	97.70 min	107.79 min	97.31 min
Maximum	N/A	97.70 min	153.48 min	103.87 min
Anomaly Analysis				
Minimum	N/A		427.29 min	340.36 min
Mean	240.35 min	Not Applicable	432.22 min	342.38 min
Maximum	N/A		439.50 min	344.70 min

Table B–2. Comparative Response Times (workload = 100 + 10 scenes/day)

Comparative Response Times by Type of Processing				
Response Time	**Service Time**	**Nominal Nom=100/Day Reproc=10/Day AAS = 0/Day**	**Batch Nom=100/Day Reproc=10/Day AAS = 6/Day**	**Uniform Nom=100/Day Reproc=10/Da y AAS = 6/Day**
Nominal				
Minimum	N/A	103.73 min	103.73 min	103.73 min
Mean	103.73 min	103.73 min	121.22 min	105.82 min
Maximum	N/A	103.73 min	211.74 min	123.66 min
Anomaly Analysis				
Minimum	N/A		568.22 min	367.45 min
Mean	362.19 min	Not Applicable	584.11 min	368.65 min
Maximum	N/A		596.38 min	370.03 min

The response time for the anomaly analysis processing includes benchmark run, transferring benchmark results to the AAS workstation, quick visual assessment of the benchmark results, diagnostic run, transferring diagnostic results to the AAS workstation, visual assessment of the diagnostic results, reprocessing the data with corrected parameters, and transferring the LPGS product.

From Table B–1 and Table B-2, it can be seen that starting all anomaly scenes at the beginning of simultaneously introduces significant queuing effects into the system. The mean and maximum response times are significantly higher than for either the Nominal or Uniform scenarios. In particular, the mean response time for the anomaly processing thread is nearly 90 minutes greater for the Batch case than for the Uniform case at 25 scenes per day, and the response time increases by 60 percent when the workload is increased to 100 scenes per day. anomaly analysis is not executed for the Nominal case.

Table B–1 and Table B-2 indicate that spreading anomaly analysis processing over three shifts has the least impact on both nominal and anomaly analysis turnaround times.

Another strategy to balance the workload during the prime shift is to reduce or curtail nominal processing of new scenes until anomaly analysis is complete. This action increases the resource loads in the remaining two shifts. Although explicit simulation runs have not been made for this scenario, such a strategy would make optimum use of available LPGS labor on a daily basis because of the intensive human interaction required for anomaly analysis.

Table B–3 and Table B-4 display the capability and offered load for each of the critical LPGS hardware devices. The offered load of a device is the average amount of work (instructions to execution of data to transfer) presented to a device relative to the device's capacity. The offered load ignores contention and synchronization. The offered load is the long-run utilization of the LPGS devices.

Table B–3. Offered Load (workload = 25 + 3 scenes/day)

Hardware Item	Capability	Nominal Nom=25/Day Reproc=3/Day AAS = 0/Day	Batch Nom=25/Day Reproc=3/Day AAS = 3/Day	Uniform Nom=25/Day Reproc=3/Day AAS = 3/Day
SGI Origin 2000	4 CPUs @ 90.675 MIPS	42.50%	51.91%	51.91%
RAID (with visual assessment included)	1 RAID group @ 70.0 MBPS	10.55%	13.14%	13.14%
FDDI	1 FDDI @ 60 Mbps	7.33%	13.51%	13.51%

Table B–4. Offered Load (workload = 100 + 10 scenes/day)

Hardware Item	Capability	Nominal Nom=100/Day Reproc=10/Da y AAS = 0/Day	Batch Nom=100/Day Reproc=10/Da y AAS = 6/Day	Uniform Nom=100/Day Reproc=10/Da y AAS = 6/Day
SGI Origin 2000	16 CPUs @ 83.85 MIPS	45.13%	50.23%	50.23%
RAID (with visual assessment included)	4 RAID groups @ 70.0 MBPS	10.36%	11.66%	11.66%
FDDI	1 FDDI @ 60 Mbps	28.80%	41.15%	41.15%

The processor speed of the SGI Origin 2000 was determined as

SGI_Speed_MIPS = (195.0 MHz / 2.0 cycles per instruction) x (1- derating factor)

where the derating factor for a four-processor configuration is seven percent and the derating factor for 16 processors is 14.00 percent.

Some of the entries in Table B–3 and Table B-4 vary slightly from the spreadsheet results because of the way in which service times were calculated.

The offered load of the Batch scenario is identical to that of the Uniform scenario because the overall arrival rate of the two is the same, even though the manner in which anomaly analysis processing is scheduled is very different.

Although the FDDI introduces only a small delay when LPGS processes 25 scenes per day, the FDDI delay increases substantially when there are 100 scenes per day. The results are shown in Table B-5. The waiting time is the delay waiting for service. Note that starting all anomaly analysis processing simultaneously significantly increases the waiting time. Other critical LPGS resources do not show such significant delays because of their low utilization.

Table B–5. FDDI Waiting Time

Scenes/ Case	Nominal	Batch	Uniform
25 + 3	0.00 min	2.75 min	0.22 min
100 + 10	0.00 min	10.58 min	1.02 min

B.3.2 Dynamic Results – 25 Scenes per Day

This scenario consists of three cases: Nominal, Batch, and Uniform. In all three cases, a total of 28 scenes are processed (including three scenes are reprocessed). In the Nominal case, all 28 scenes receive nominal processing, but no anomaly analysis processing. In the Batch case, 28 scenes receive only nominal processing, and three cases also receive anomaly analysis processing. Anomaly analysis processing is done once per day in the Batch case with all anomaly cases are started simultaneously. In the Uniform case, 28 scenes receive only nominal processing, three cases also receive anomaly analysis processing, and the anomaly cases start at equal intervals throughput the day. There are four CPUs derated by seven percent, and there is a single RAID group.

B.3.2.1 Nominal Case

In the Nominal case, none of the resources become saturated, and the response times are virtually constant with the response time approximately equal to the total service time.

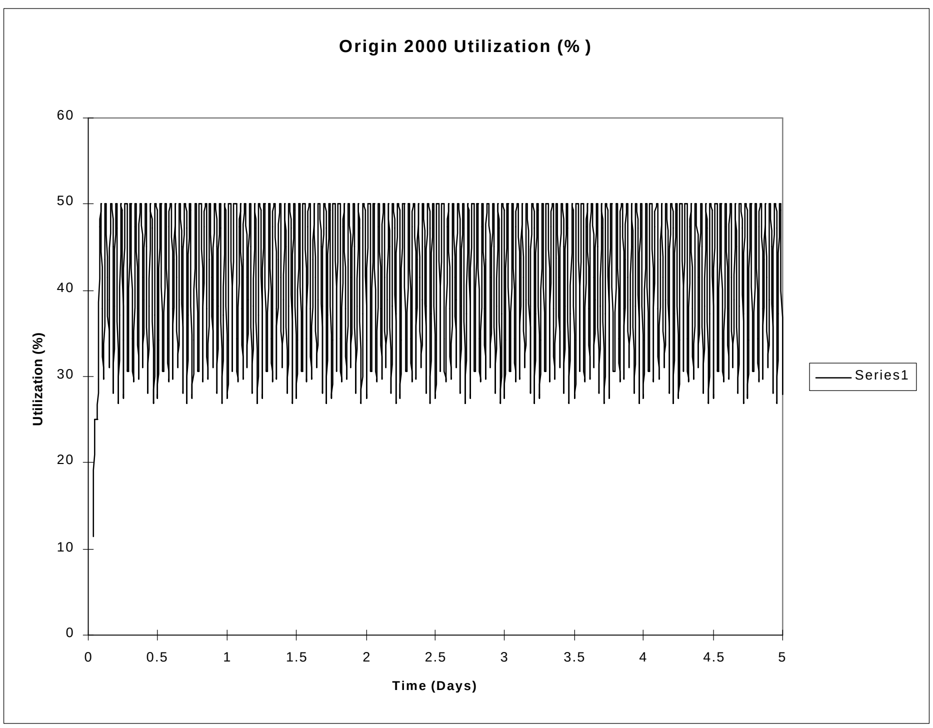

Figure B-3. Origin 2000 Utilization (Nominal Case)
(Workload = 25 + 3 Scenes per Day)

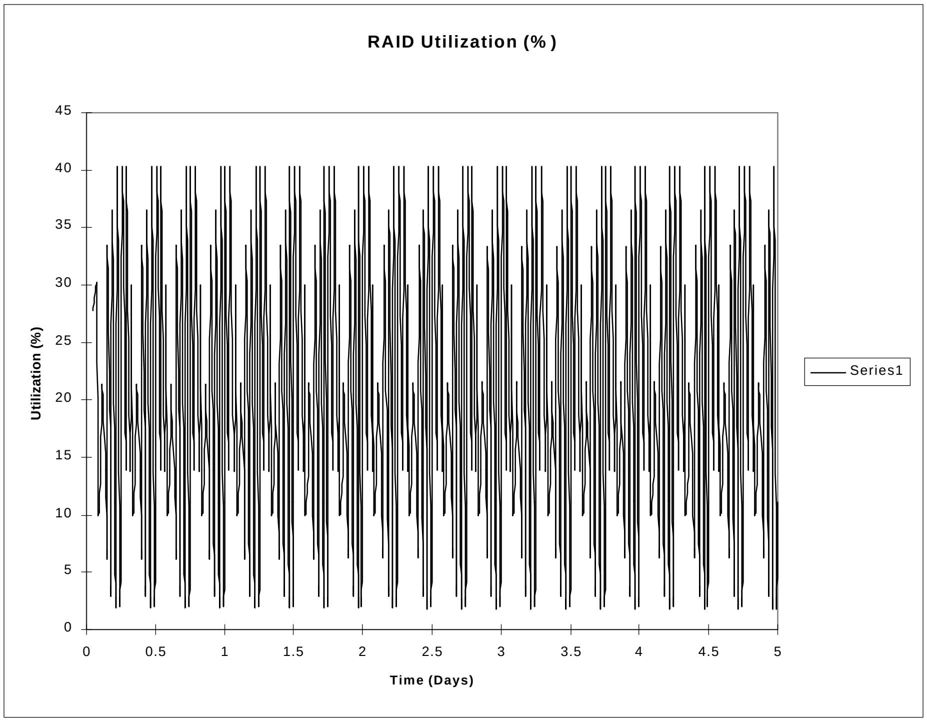

Figure B-4. RAID Utilization (Nominal Case)
(Workload = 25 + 3 Scenes per Day)

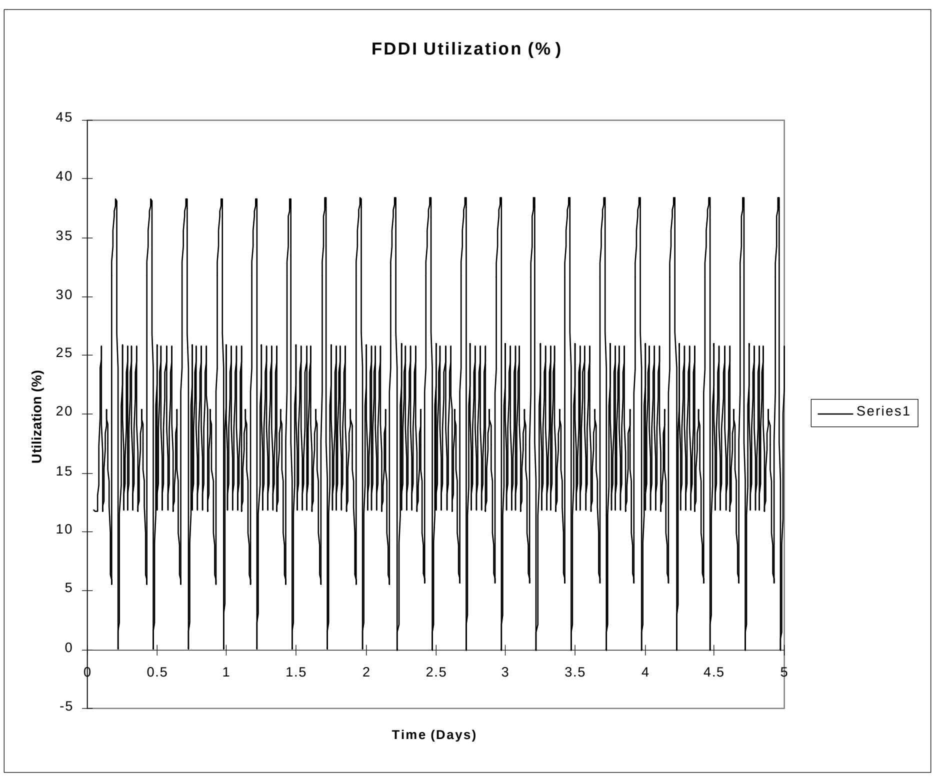

Figure B-5. FDDI Utilization (Nominal Case)
(Workload = 25 + 3 Scenes per Day)

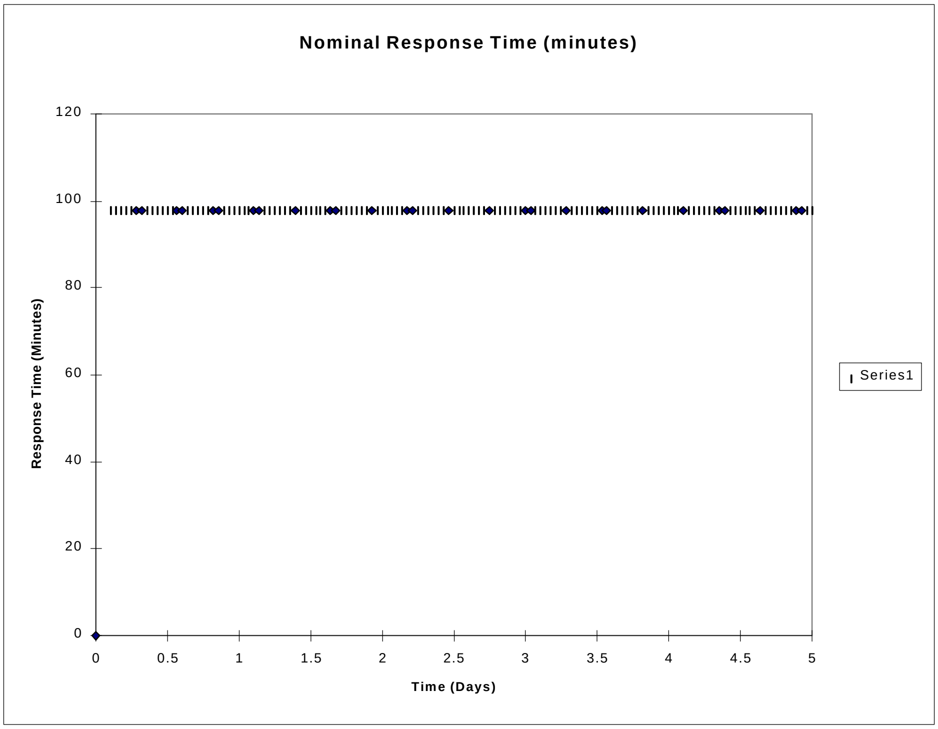

Figure B-6. Nominal Processing Response Times (Nominal Case)
(Workload = 25 + 3 Scenes per Day)

B.3.2.2 Batch Case

In the batch case, there are three anomaly scenes which are started simultaneously. This scheduling policy clearly causes both the Origin 2000 processor and the FDDI to become saturated. There is also a clear effect on the anomaly analysis response times due to network queuing.

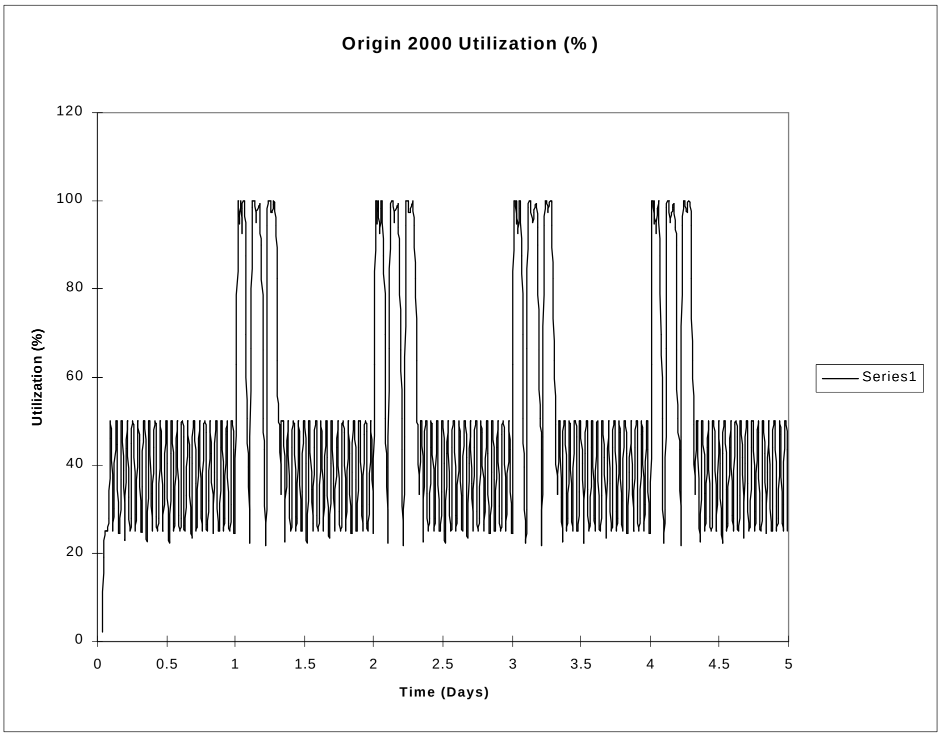

B-7. Origin 2000 Utilization (Batch Case)
(Workload = 25 + 3 Scenes per Day)

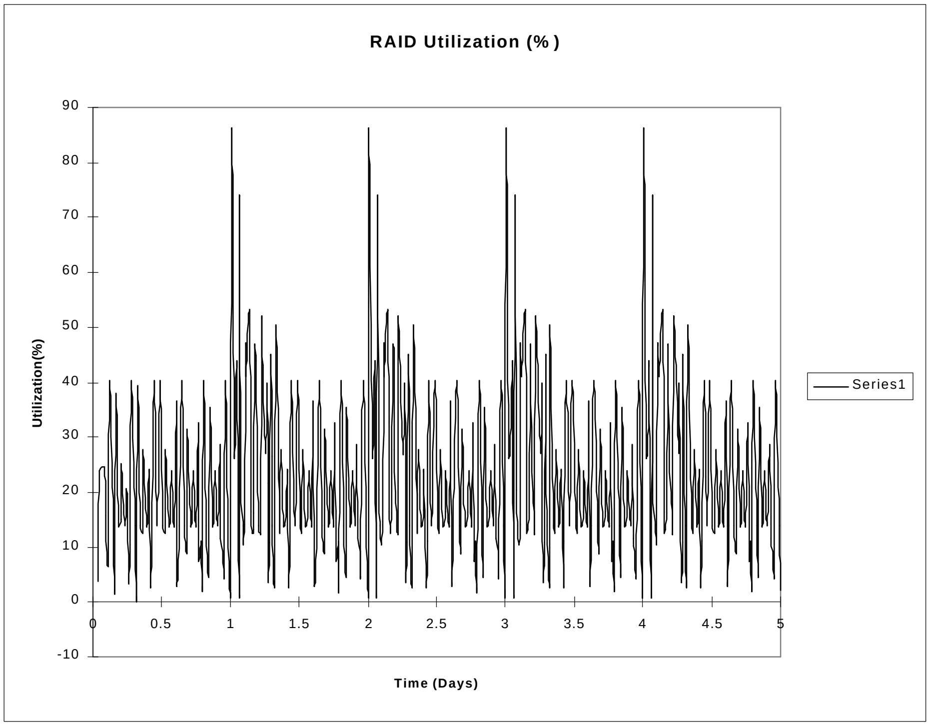

Figure B-8. RAID Utilization (Batch Case)
(Workload = 25 + 3 Scenes per Day)

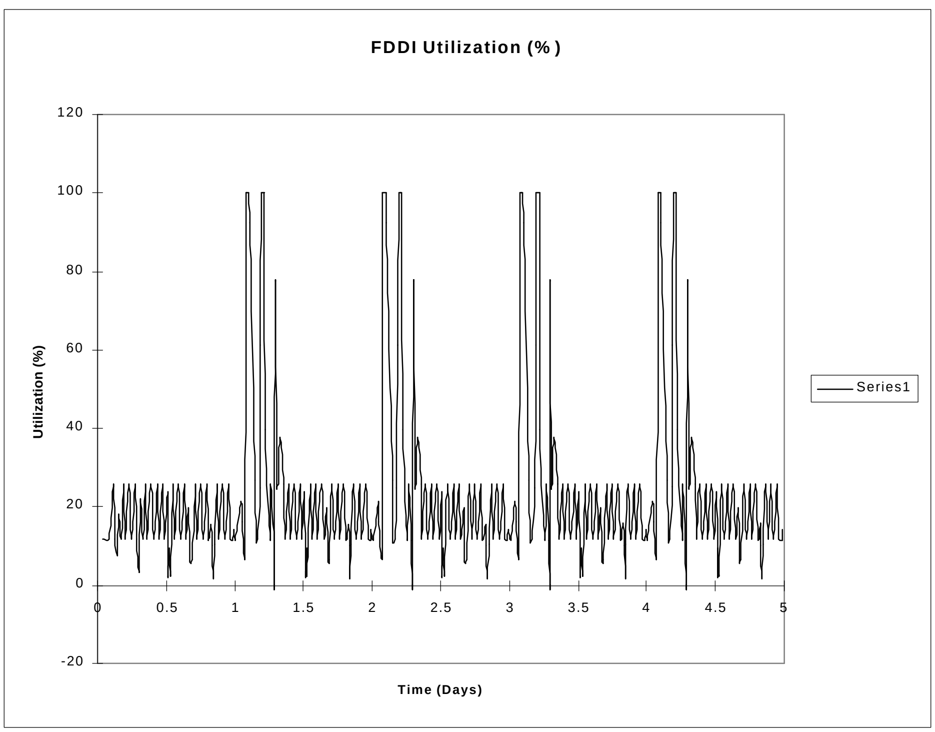

Figure B-9. FDDI Utilization (Batch Case)
(Workload = 25 + 3 Scenes per Day)

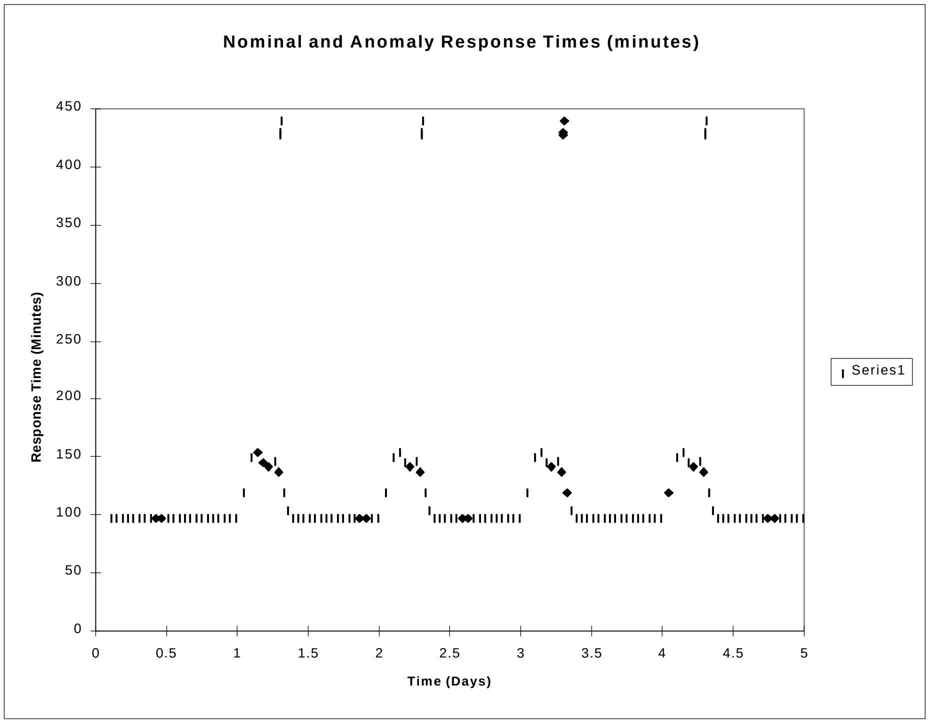

Figure B-10. Nominal and Anomaly Response Times (Batch Case) (Workload = 25 + 3 Scenes per Day)

B.3.2.3 Uniform Case

In this scenario, there are 28 WS scenes per day, and three anomaly scenes. The three anomaly scenes are processed at equal intervals over the course of a day. The FDDI still becomes saturated for brief periods, but there is only a small effect on anomaly analysis response times.

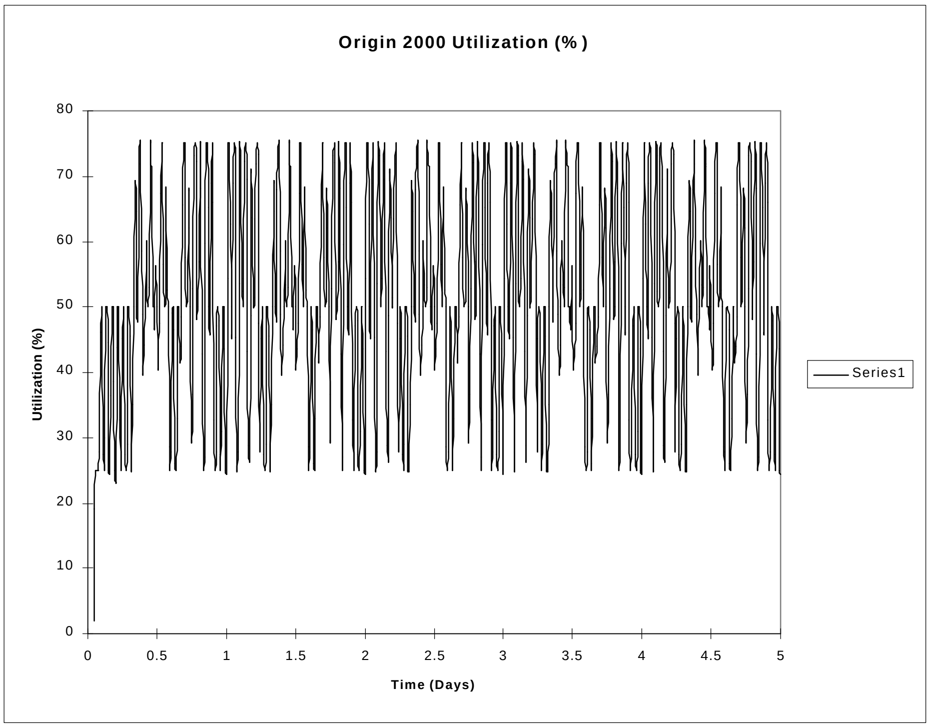

Figure B-11. Origin 2000 Utilization (Uniform Case) (Workload = 25 + 3 Scenes per Day)

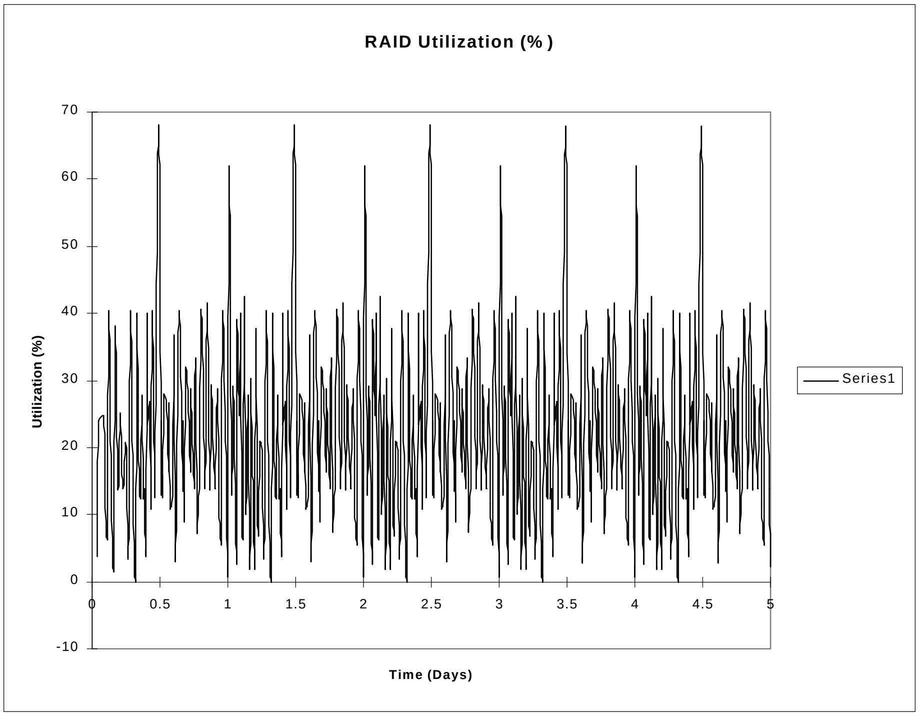

Figure B-12. RAID Utilization (Uniform Case)
(Workload = 25 + 3 Scenes per Day)

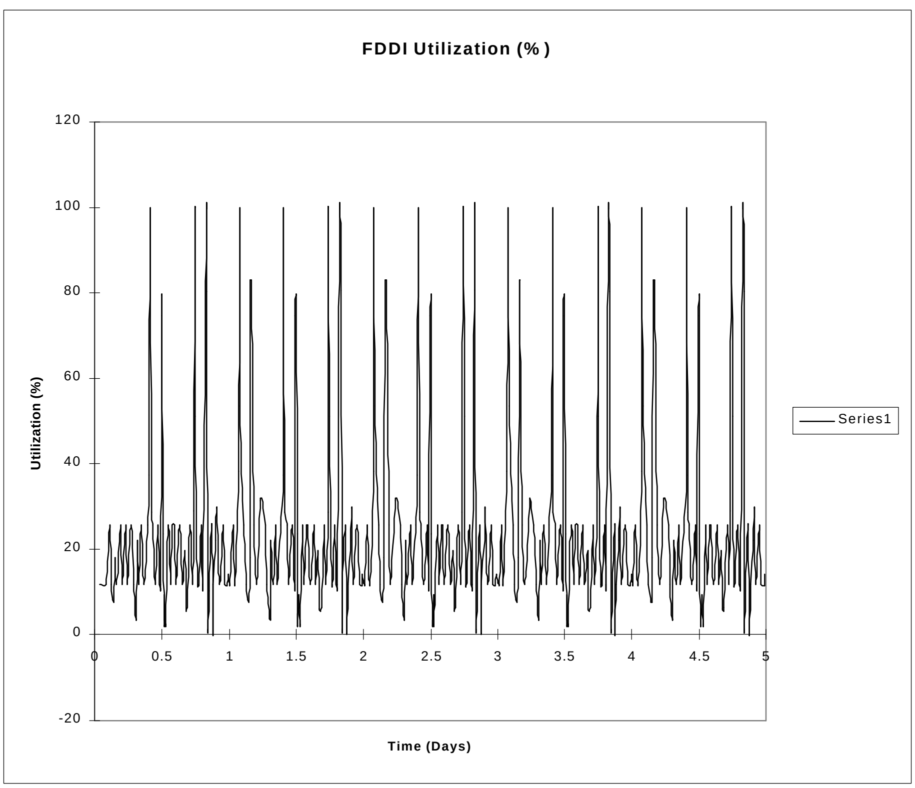

Figure B-13. FDDI Utilization (Uniform Case)
(Workload = 25 + 3 Scenes per Day)

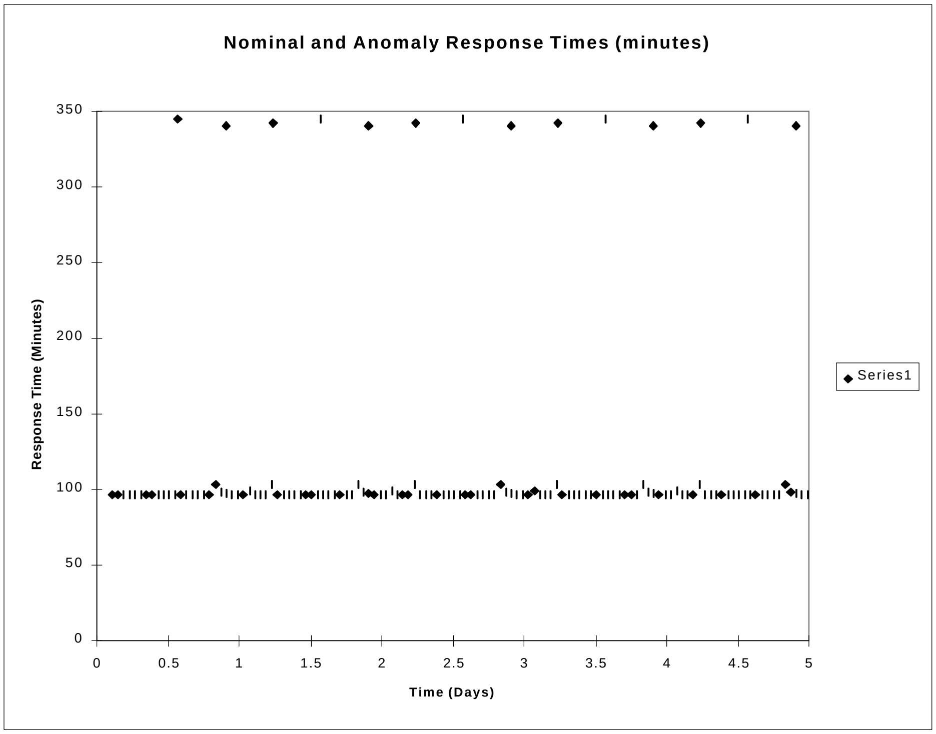

Figure B-14. Nominal and Anomaly Response Times (Uniform Case) (Workload = 25 + 3 Scenes per Day)

B.3.3 Dynamic Results – 100 Scenes per Day

This scenario consists of three cases: Nominal, Batch, and Uniform. In all three cases, a total of 110 scenes are processed (including ten scenes are reprocessed). In the Nominal case, all 110 scenes receive nominal processing, but no anomaly analysis processing. In the Batch case, 110 scenes receive only nominal processing, and ten cases also receive anomaly analysis processing. Anomaly analysis processing is done once per day in the Batch case with all anomaly cases started simultaneously. In the Uniform case, 110 scenes receive only nominal processing, ten cases also receive anomaly processing, and the anomaly cases start at equal intervals throughput the day. There are 16 CPUs derated by 14 percent, and there are four RAID groups. There is only a single FDDI.

B.3.3.1 – Nominal Case

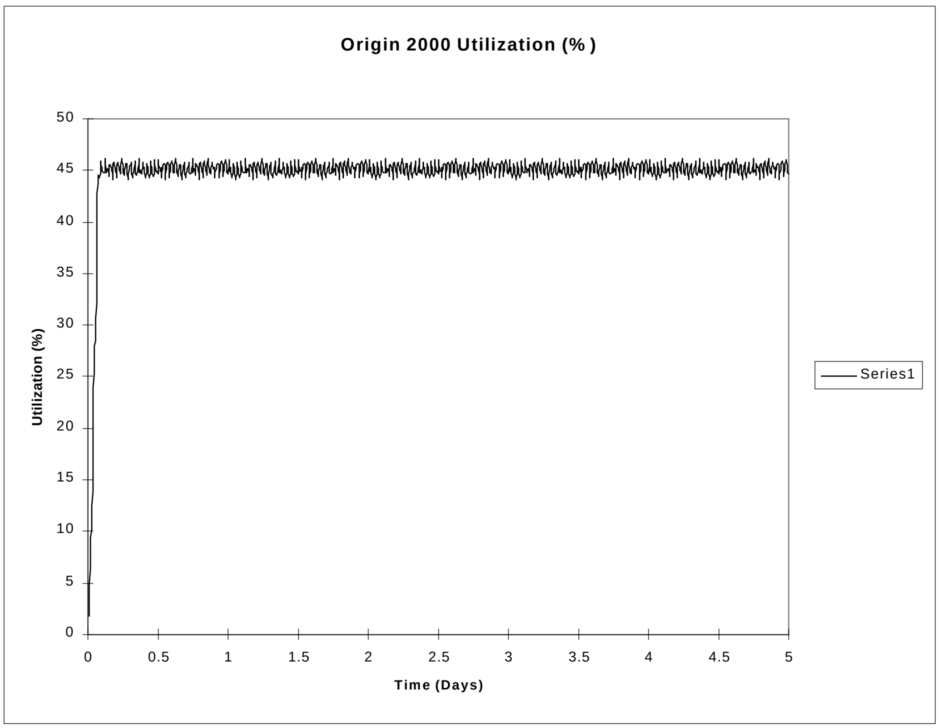

Figure B-15. Origin 2000 Utilization (Nominal Case)
(Workload = 100 + 10 Scenes per Day)

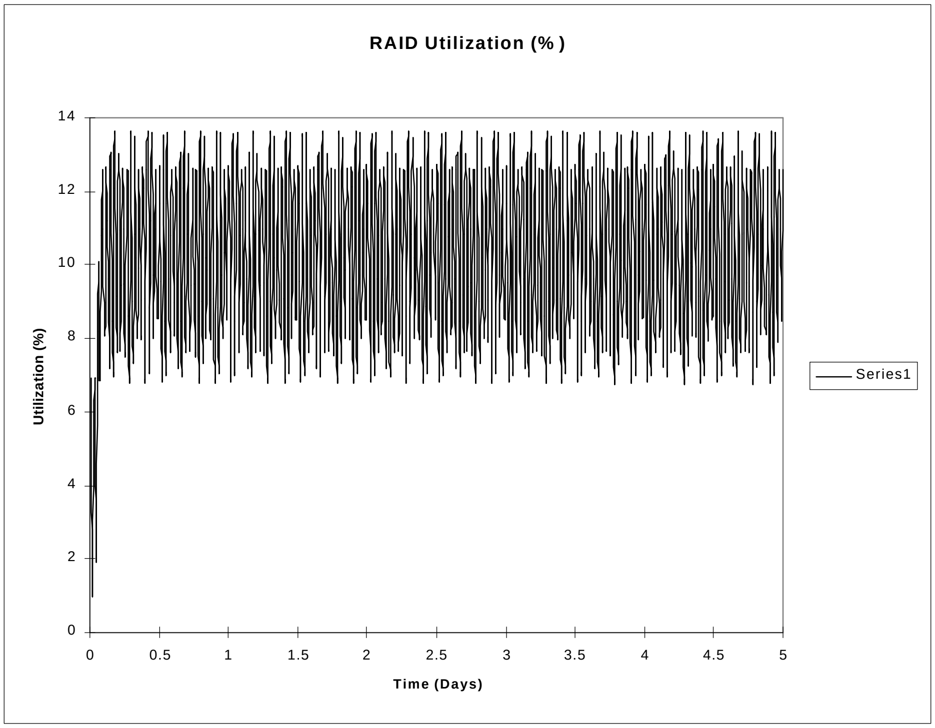

Figure B-16. RAID Utilization (Nominal Case)
(Workload = 100 + 10 Scenes per Day)

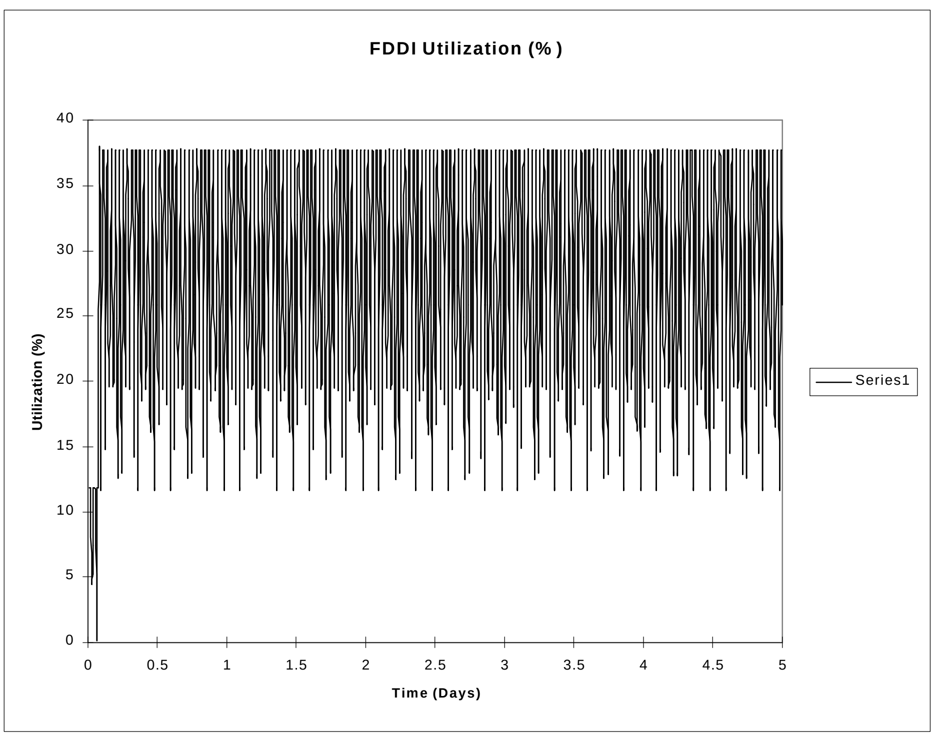

Figure B-17. FDDI Utilization (Nominal Case)
(Workload = 100 + 10 Scenes per Day)

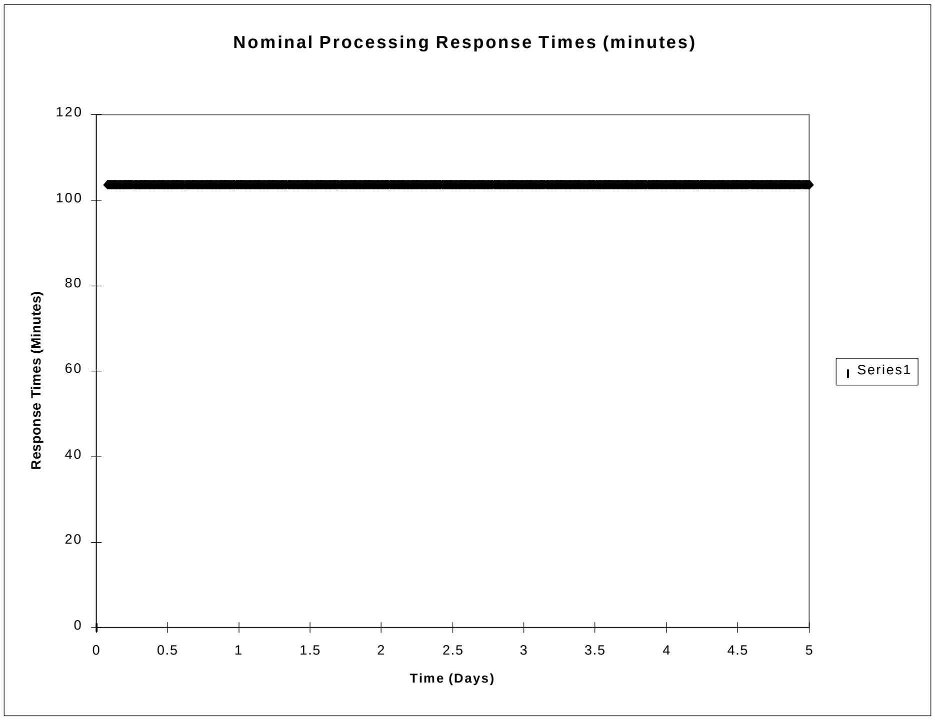

Figure B-18. Nominal and Anomaly Response Times (Nominal Case)
(Workload = 100 + 10 Scenes per Day)

B.3.3.2 Batch Case

As in the previous Batch case, there are ten anomaly scenes which are started simultaneously. This scheduling policy clearly causes both the Origin 2000 processor and the FDDI to become saturated. There is also a clear effect on the nominal response times due to network queuing primarily on the FDDI.

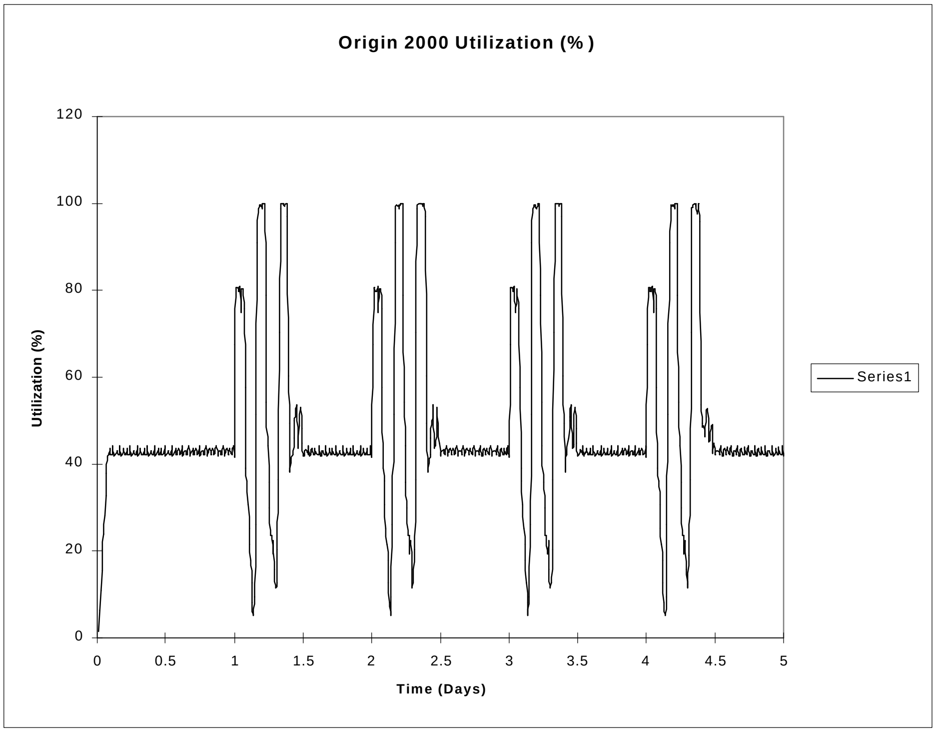

Figure B-19. Origin 2000 Utilization (Batch Case)
(Workload = 100 + 10 Scenes per Day)

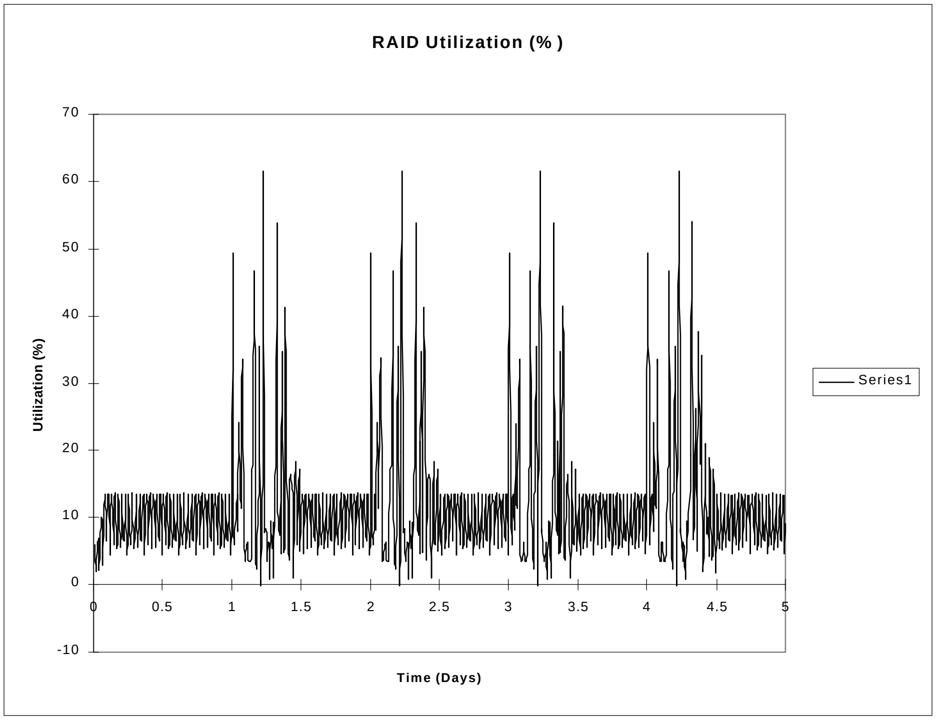

Figure B-20. RAID Utilization (Batch Case)
(Workload = 100 + 10 Scenes per Day)

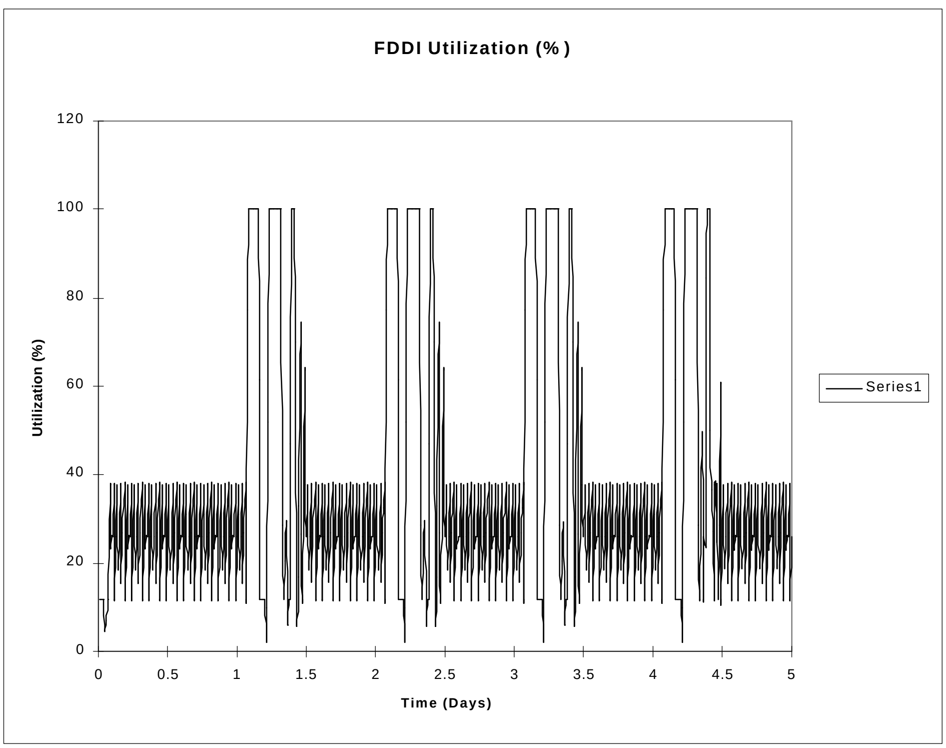

Figure B-21. FDDI Utilization (Batch Case)
(Workload = 100 + 10 Scenes per Day)

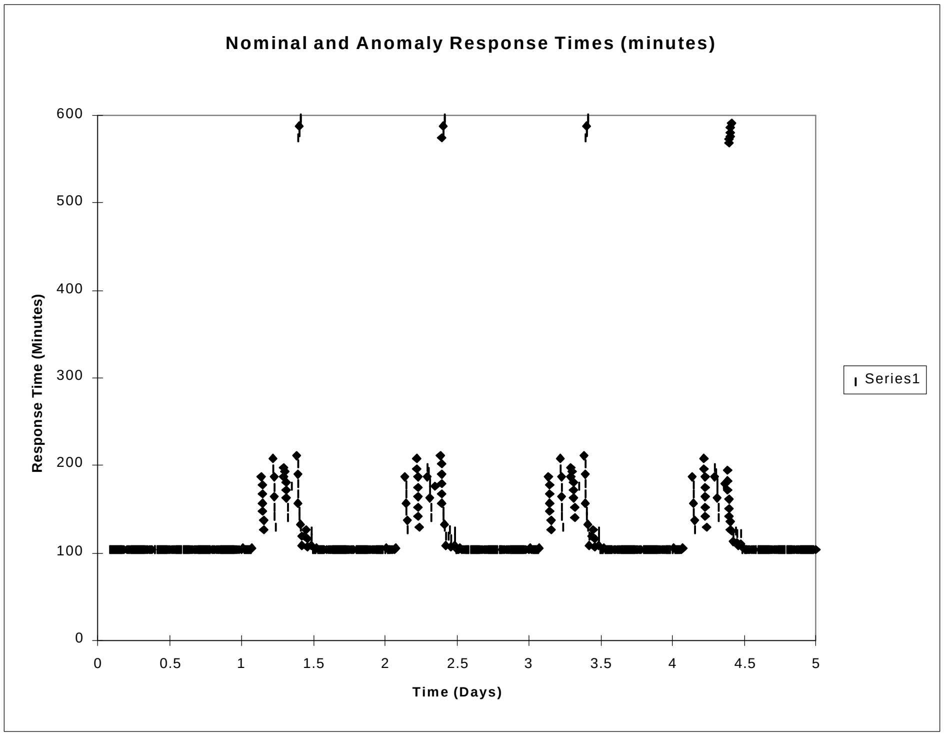

Figure B-22. Nominal and Anomaly Response Times (Batch Case) (Workload = 100 + 10 Scenes per Day)

B.3.3.3 - Uniform Case

In this scenario, there are 110 WS scenes per day, and ten anomaly scenes. The anomaly scenes are processed at equal intervals over the course of a day. Both the CPU and the FDDI become saturated for brief periods, but there is only a small effect on nominal response times, and virtually no effect on the anomaly analysis response times.

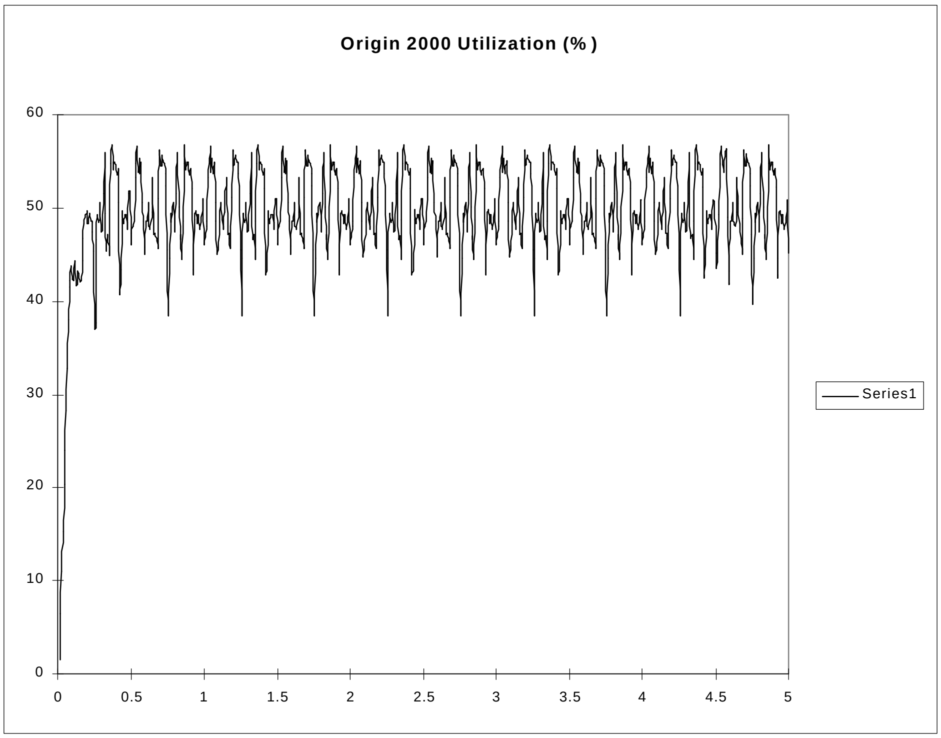

Figure B-23. Origin 2000 Utilization (Uniform Case) (Workload = 100 + 10 Scenes per Day)

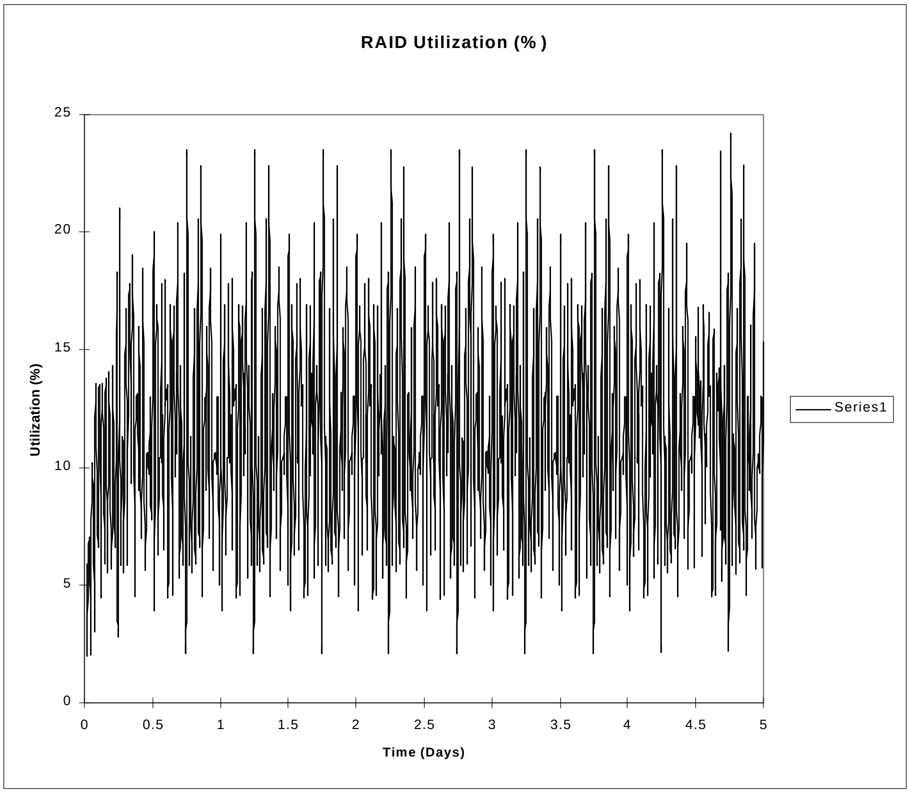

Figure B-24. RAID Utilization (Uniform Case)
(Workload = 100 + 10 Scenes per Day)

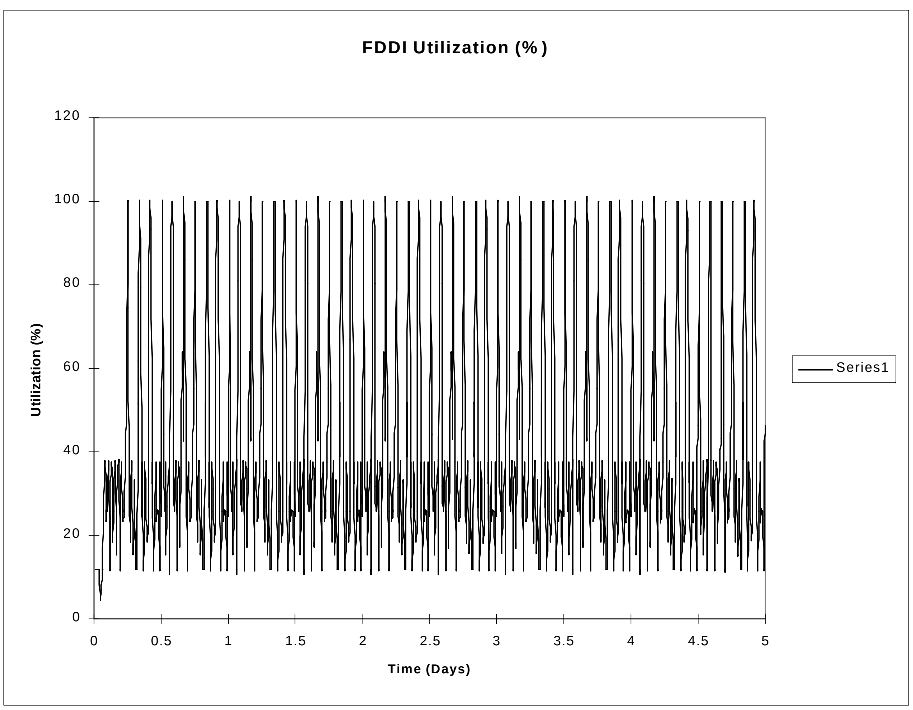

B-25. FDDI Utilization (Uniform Case)
(Workload = 100 + 10 Scenes per Day)

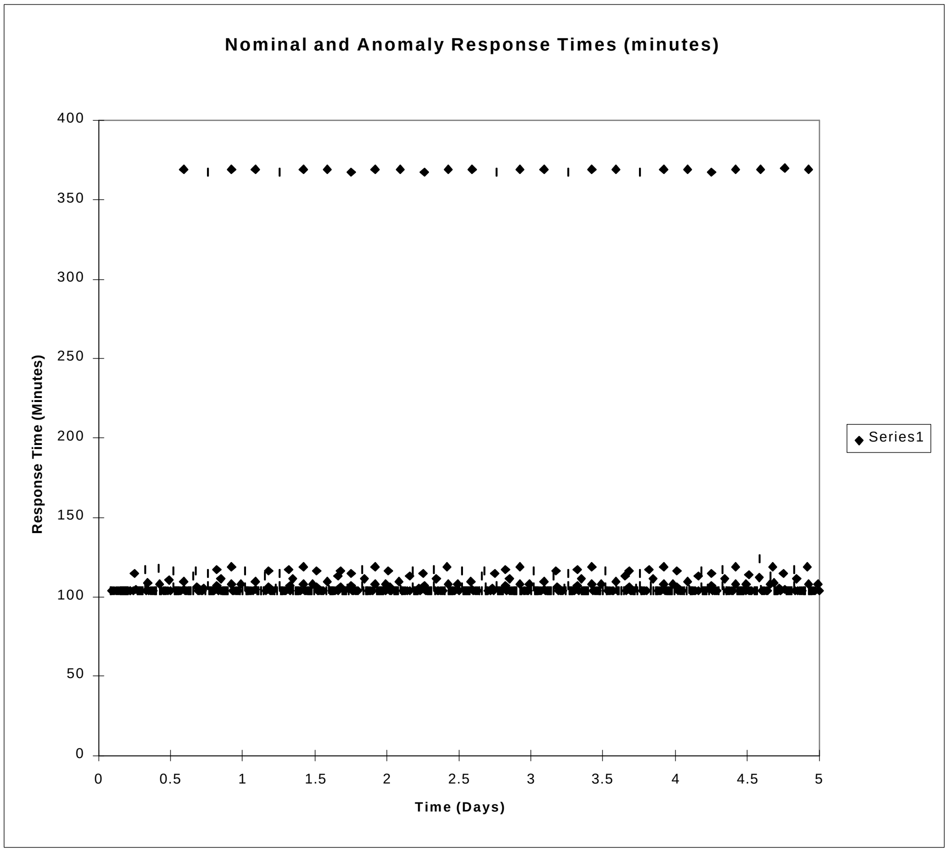

Figure B-26. Nominal and Anomaly Response Times (Uniform Case)
(Workload = 100 + 10 Scenes per Day)

Appendix C. Processing Scenarios

Even though the compiler can optimize the code for multiprocessor systems, to fully take advantage of multiple CPUs, the application software needs to be designed to allow parallel processing on the data. Depending on the degree of parallel processing of the application software, the following different processing scenarios could arise. A four-CPU configuration with a 7-percent CPU performance degradation will be assumed in the following discussions.

C.1 Scenario 1: Sequential Processing of All Bands, One Band After Another

This scenario assumes that the application software does not provide the capability for simultaneously processing different bands of the same work order. Therefore, all bands from the same work order can only be processed sequentially. To take advantage of multiple CPUs, multiple work orders need to be processed simultaneously, one on each CPU. This scenario will not require any synchronization and can be supported by the current software design.

C.2 Scenario 2: Process Multiple Bands Simultaneously

This scenario assumes that parallel processing can be done on different bands of the same work order on different CPUs simultaneously. Because Band 8 takes almost four times longer to process, the bottleneck for processing a WRS scene is in the processing of this band. Table C-1 shows the amount of time it takes to process a single band of data. Note that the Ingest Data, Format Product, and Transfer Product are for the entire scene; Table C–1 shows no breakdown for these three processes.

This scenario will require some synchronization before the next processes, such as L1G Processing and Format Product, can proceed. However, the current design of the radiometric processing can not support this scenario without design changes.

C.3 Scenario 3: Parallel Processing Within a Band

This scenario will provide a maximum flexibility in processing the data. But it will require that the software be able to break the data and work to allow parallel processing (or multithreading) on data within a band. It allows data and work for a band being processed by multiple CPUs simultaneously. There will be additional overhead due to the necessary synchronization of processes. Total service time may increase slightly due to the overhead but the total wall clock time to process a band of data can be significantly reduced. Another advantage of this scenario is that the memory requirements can be significantly reduced. 1.73 GB of memory for the entire system (instead of 1.73 GB per CPU) for the geometric processing will be sufficient if all four CPUs are processing the same band of data simultaneously. Additional memory will allow processing of multiple bands or multiple work orders simultaneously without the data being unnecessarily swapped in and out.

The current geometric processing software design supports this scenario. Furthermore it does not require the entire band of data to be resident in the memory. It allows processing on smaller amounts of data within a band (e.g., 50 scans of data).

The current design of the radiometric processing can not support this scenario without design changes.

Table C-1. Nominal Processing Time by Band for One WRS Scene

Level 1 Processing HWCI (Normal Processing) - Process/Data Transfer Time (by band for 1 CPU) CPU performance degradation factor=7%							
in minutes	**# of Occur-rences**	**Ingest Data**	**L1R Process**	**L1G Process**	**Format Product**	**Transfer Product**	**Total for All Bands**
CPU Time							
Band 1/2/3/4/5/7	6		4.24	3.10			44.04
Band 6	1		1.17	2.28			3.45
Band 8	1		25.66	12.40			38.06
All Bands	1	0.55			0.64	0.77	1.96
Data Transfer Time							
Band 1/2/3/4/5/7	6		0.27	0.13			2.40
Band 6	1		0.08	0.04			0.12
Band 8	1		1.02	0.51			1.53
All Bands	1	1.57			0.83	2.86	5.26
Total		2.12	54.99	34.61	1.47	3.63	

C.4 Discussion

Table C-2 shows the minimum wall clock times for different combinations of the above scenarios. These minimum wall clock times are derived from the data in Table C-1. Wall clock times corresponding to the scenarios supported by the current software design are shown in ***Italic Bold***. Please note that the wall clock time estimated in Table C-2 is assuming that only one work order is being processed in a four-CPU configuration. The wall clock time will significantly increase if any other jobs in addition to the work order are running on the system.

In the actual situation with multiple CPUs, the operating system will schedule many jobs (either of the same work orders or different work orders) simultaneously. Each job will get a slice of total CPU time. As the number of jobs increases, the wall clock time to complete each job increases. The number of work orders/bands that can be processed simultaneously is determined by the amount of memory available. Insufficient memory will result in the data being swapped in and out unnecessarily, which is very inefficient for the size of the data processed by the LPGS.

Table C-2. Minimum Wall Clock Times for Processing One WRS Scene With Different Scenarios in a Four-CPU Configuration

Radiometric Processing	Geometric Processing		
	Scenario 1	Scenario 2	Scenario 3
	Sequential processing, one band after another	Processing multiple bands simultaneously	Parallel processing within a band
Scenario 1 Sequential processing, one band after another	***96.61 minutes***	***75.12 minutes (1)***	***71.86 minutes (2)***
Scenario 2 Processing multiple bands simultaneously	68.51 minutes (3)	46.81 minutes (4)	37.51 minutes (5)
Scenario 3 Parallel processing within a band	57.62 minutes (6)	35.92 minutes (7)	32.66 minutes (8)

Notes:

(1) (2.12+54.99+12.40+0.51+1.47+3.63) (bottleneck: radiometric processing, geometric processing of Band 8 data)

(2) (2.12+54.99+(3.10*6+2.28+12.40)/4+(0.13*6+0.04+0.51)+1.47+3.63) (bottleneck: radiometric processing)

(3) (2.12+25.66+1.02+34.61+1.47+3.63) (bottleneck: radiometric processing of Band 8 data and geometric processing)

(4) (2.12+25.66+1.02+12.40+0.51+1.47+3.63) (bottleneck: processing of Band 8 data)

(5) (2.12+25.66+1.02+12.40/4+0.51+1.47+3.63) (bottleneck: radiometric processing of Band 8 data)

(6) (2.12+(4.24*6+1.17+25.66)/4+(0.27*6+0.08+1.02)+34.61+1.47+3.63) (bottleneck: geometric processing)

(7) (2.12+(4.24*6+1.17+25.66)/4+(0.27*6+0.08+1.02)+12.40+0.51+1.47+3.63) (bottleneck: geometric processing of Band 8 data)

(8) (2.12+(44.04+3.45+38.06)/4+(2.40+0.12+1.53)+1.47+3.63)

Acronyms and Abbreviations

AA	anomaly analysis
AAS	Anomaly Analysis Subsystem
AST	Advanced System Technology, Inc.
CPU	Central Processing Unit
DAAC	Distributed Active Archive Center
ECS	EOS Core System
EDC	EROS Data Center
EGS	ESDIS Ground System
ENVI	Environment for Visualizing Images
EOS	Earth Observing System
EOSDIS	EOS Data and Information System
EROS	Earth Resources Observation System
ESDIS	Earth Science Data and Information System
ETM+	Enhanced Thematic Mapper Plus
F&PRS	functional and performance requirements specification
FDDI	fiber-optic data distribution interface
FIFO	first in-first out
GB	gigabyte
HWCI	hardware configuration item
IAS	Image Assessment System
IDL	Interactive Data Language
I/O	input/output
L0R	Level 0 radiometrically corrected
L1	Level 1
L1G	Level 1 geometrically corrected
L1R	Level 1 radiometrically corrected
LPGS	Level 1 Product Generation System
MB	megabyte
MFLOPS	Million Floating Point Operations

MHz	megahertz
MSCD	mirror scan correction data
PCD	payload correction data
QA	quality assessment
QASE RT	Quantitative Case for Reliability and Timing
RAID	redundant arrays of independent disks
SCSI	small computer system interface
SGI	Silicon Graphics, Inc.
WRS	Worldwide Reference System